

BACKGROUND OF THE EDITION

CROSS-RELATED APPLICATION

This application takes priority from Provisional Patent Application Serial No. 60/261,932 filed January 15, 2001.

FIELD OF THE INVENTION

The present invention relates to certain improvements in digital video recording.

BACKGROUND ART

Compressed video frames are variable in size depending on the complexity of the picture. Therefore, the recorder must maintain a directory of video frame disk addresses to support cueing (jumping to an arbitrary frame), random access playback (switching seamlessly without hesitation or interruption among material recorded at different physical locations on disk), and “feature” playback (fast forward and reverse scan).

Typically, this frame directory is kept on disk. With this approach, when a user submits a cue command, the recorder must read the directory first to determine the disk address, then read the frame from disk. This latency results in a visible hesitation or requires that the cue command be received some period in advance of the required cue. There is a need to overcome such latency.

Conceptually, recording and playing motion video to and from hard disk is simply performing the conversion from its analog (or uncompressed digital) format to a continuous compressed data stream and managing this continuous stream of data as a series of "writes" or "reads to" or "reads from" hard disk.

1 The hard disk is the slowest component in the system. To achieve the
2 performance threshold that broadcast quality playback requires, it is necessary to
3 interface to the hard disk at maximum efficiency. The data transfer method must
4 be tailored to the hard disk's characteristics. With modern hard disks, the highest
5 performance interface is achieved with burst data transfers of a size that matches
6 the disk's cache memory. There is a need to implement burst data transfers
7 appropriate for video data at maximum efficiency.

8
9 The user configures a recorder to record video at a specified compression
10 ratio, for example, 4:1. The disk may not be able to sustain this specified data
11 rate when the user has installed a low performance disk or if the image data
12 frequency changes from simple pictures (black frames, for example) to complex
13 ones. For example, one may watch MTV for a few minutes for excellent
14 examples of video that stress JPEG compression rate-averaging algorithms. If
15 the disk falls far enough behind, the result is catastrophic: The record stream
16 skips frames, losing them irretrievably. There is a need to control compression
17 so that such catastrophic results are averted.

18
19 Certain applications benefit from the ability to record continuously. In
20 surveillance, for example, it is useful to record until an alarm occurs, then to stop
21 some time after the alarm. The recorder has captured time before, during, and
22 after the alarm, which can be useful in determining the cause of the alarm. There
23 is a need to provide support for such a loop recording mode.

SUMMARY OF THE INVENTION

To overcome directory-based latency, the recorder of the present invention reads the disk addresses of all recorded frames into memory on startup and maintains this index table as material is recorded and deleted. The in-memory index table eliminates a disk access before loading the required frame. The benefit is instantaneous cueing (latency of only one frame time), on-the-fly changes to the playback sequence, and generally quick response to user motion control. The cost is providing adequate memory to contain the address of every frame (at 108,000 frames per hour). The invention provides adequate indexable space to support random access within eight hours of recorded video, expandable to more than forty hours with an "extended time module" daughter card.

The invention's memory controller, implemented in the preferred embodiment in a Xilinx FPGA, accommodates the disk's data transfer requirements by performing this burst DMA transfer in programmable size data blocks. A memory controller makes extensive use of the "extended data out" (EDO) feature of dynamic random-access-memory to enhance data transfer speeds.

Further, the memory controller must arbitrate access among several additional diverse devices which share the same memory. The controller must also support the unique data transfer characteristics of these other devices without compromising the burst DMA timing required by the hard disk interface.

- 1) Compressed video DMA must move continuously, with possible extreme peaks in the data rate. Average data rate is at least 5 megabytes per second for broadcast quality;
- 2) The graphics overlay requires memory accesses at fixed intervals in sync with the video raster, fixed data rate of about 5 megabytes per second;

- 1 3) The uncompressed video channel, when used, requires fixed
- 2 timing at 27 megabytes per second;
- 3 4) The CPU has latency restrictions and must complete memory
- 4 access within stringent limits. The amount of data transferred
- 5 by the CPU includes digital audio, running at 192 kilobytes
- 6 per second;
- 7 5) The memory requires refresh cycles with a maximum latency
- 8 restriction.

9

10 The invention's firmware supports a "loop record" mode where the recorder

11 first fills available space, then continues recording by overwriting the oldest

12 video. To continue this process indefinitely is difficult because the video frames

13 vary in size and the index table, like the video frames themselves, must be

14 maintained in a circular manner. Further, the recorder must be able to stop at

15 any moment and be able to play back the video as expected (oldest to newest)

16 regardless of how it is actually stored on disk.

17

18 Another application that takes advantage of loop recording and adds another

19 level of complexity is "time delay", where the device continuously records a

20 stream of video and plays the same stream back some fixed interval of time later.

21 The invention's firmware supports this application where two circuit boards share

22 a common disk array. One board is dedicated to recording, the other board to

23 playback. The two circuit boards communicate over a serial link. The recording

24 board passes the information required for playback to the playback board. This

25 configuration can also be used for "instant replay" type applications because the

26 playback board can randomly access and play back the majority of the video

27 frames stored on disk by the recording board.

BRIEF DESCRIPTION OF THE DRAWINGS

3 The aforementioned objects and advantages of the present invention, as well
4 as additional objects and advantages thereof, will be more fully understood
5 hereinafter as a result of a detailed description of a preferred embodiment when
6 taken in conjunction with the following drawings in which:

8 FIG. 1 is a block diagram of the architecture of the video recorder of a
9 preferred embodiment of the invention:

11 FIG. 2 is a block diagram representation of the relation between hard disks
12 and index tables used in the preferred embodiment:

4 FIG. 3 is a block diagram representation of the relation between individual
5 video frames and index table regions:

17 FIG. 4 is a block diagram representation of the relation between available
18 records, disk record directory entries and index table regions;

20 FIG. 5 is a block diagram representation of the relation between available
21 records and index table space for use in logs records:

23 FIG. 6 is a general block diagram of an arbitration assembly for multi-port
24 access to a single memory device;

26 FIG. 7 is a block diagram similar to that of FIG. 6, but for a 4-port memory
27 system;

29 FIG. 8 is a block diagram of memory control logic employed in the system of
30 FIG. 7;

1 FIG. 9 is a logic schematic of a simplified arbiter; and

2

3 FIG. 10 is a flow chart describing a dynamic metering algorithm for controlling

4 the recorder's compression and coding rate.

5

1 **DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT**

2

3 The present invention comprises a unique high performance digital video
4 recorder having a number of novel features. The recorder's electronics are all on
5 a unitary printed circuit board, a block diagram of which is provided in FIG. 1.
6 The recorder also requires at least one hard disk drive and audio and video input
7 analog signals (from a source such as video camera or broadcast media) as well
8 as a suitable monitor for receiving output audio and video analog signals. A
9 external time code generator (i.e., VITC digital clock) is also required for
10 synchronization. Also required are various manual control devices (i.e., panel
11 controls) for mode selection as will be discussed hereinafter. The electronics of
12 the preferred embodiment comprise A-to-D and D-to-A converters, a hard disk
13 interface, a JPEG compression encoder/decoder, a multi-port DRAM and DMA
14 subsystem, a microprocessor with RS-232 and RS-422 access ports, various
15 working memory devices and bus interfaces and a 16-bit stereo digital audio
16 subsystem. The novel features of the preferred embodiment comprise use of an
17 index table for disk addresses of recorded frames, a multi-port memory controller
18 in the form of a field programmable gate array (FPGA), loop recording using dual
19 channels, and dynamic JPEG compression compensation. Each of these
20 features will now be described in sufficient detail to enable those having ordinary
21 skill in the relevant art to make and use the invention without undue
22 experimentation.

23

24 **INDEX TABLE**

25 During recording, individual frames of video are compressed and stored
26 sequentially into contiguous regions of unused disk space, called "Available
27 Records". A particular recording of a continuous sequence of video frames,
28 called a "clip", may span many Available Records across many individual disk
29 drives, up to the total amount of disk space available (see FIG. 2). The recording
30 process uses all Available Records on one particular disk drive before continuing
31 onto another disk drive. For each Available Record, the recording process is

1 provided with the disk drive unit number, the size of each addressable block on
2 that disk drive (or "block size"), and the starting and ending disk addresses of the
3 recordable area of the Available Record.

4

5 Given the starting disk address and block size, along with the size of each
6 video frame derived from the video compression process as frames are stored
7 into memory, the recording process is able to calculate the starting disk address
8 for each video frame, as well as the starting disk address for the next video frame
9 to be stored (see FIG. 2). It is these starting disk addresses that are placed into
10 the index table. For the last video frame in each Available Record, the "starting
11 address" for the "next" video frame is also placed into the index table, although
12 video is not actually stored at this disk location. (If the recording process
13 continues, this video frame would be stored at the starting address of a new
14 Available Record since the recording process is able to calculate that the video
15 frame is too large for the space left in the current Available Record.) The region
16 used to store video frames within an Available Record becomes a "Media
17 Record" (see FIG. 2). Thus for each Media Record, the corresponding region of
18 the index table consists of a sequential set of disk addresses, the number of
19 which is one more than the number of video frames in the Media Record (see
20 FIG. 3). This allows the starting disk address and number of disk blocks for each
21 frame of video on disk to be known during the playback process prior to the
22 frame being retrieved from disk.

23

24 The associated index table region for each Media Record is also stored on
25 the same disk as the Media Record as a "Control Record". (The recording
26 process insures there will be space for the Control Records left on the disk drive.)
27 While conceptually there could be a Control Record for each Media Record, in
28 the preferred embodiment the index table and the derived Control Record are
29 both continuous and contiguous for all Media Records of a single clip on a single
30 disk drive (see FIG. 2). This maintains the proper sequence for the subset of the
31 clip's video frames stored on that disk drive. (The various Control Records and

1 Media Records of a clip are associated by using a unique clip identification
2 number, or "Clip ID". The proper sequence of Control Records and Media
3 Records is maintained using a sequential "Sequence Number". These numbers
4 are also stored with each Control Record and Media Record, along with the
5 number of video frames represented or contained within the Record. It is
6 important to note that each video frame consists of two video fields, and that the
7 audio associated with a video frame may be stored with the video frame or
8 separately in "Audio Records". Since frames of audio are of a fixed size, there is
9 no need for a separate index table for audio data.)

10

11 Prior to initiating playback of any particular clip, all Control Records for that
12 clip are loaded into the index table to allow for random frame access. While this
13 must happen at least just prior to initiating playback, for increased performance,
14 the circuit card software loads all Control Records into the index table upon
15 startup, and maintains the index table as video frames are recorded and deleted,
16 and as disk drives are added and removed. When a Control Record is loaded
17 into the index table or stored after recording, the Media Records and related
18 index table regions are associated using an implementation specific identifier
19 called an index table "reference".

20

21 In the invention, the index table is implemented within a reserved area of
22 memory and managed by an implementation specific set of software routines,
23 which include index table memory management. The software interface to the
24 index table includes:

25

26 1. Initializing the index table for use upon startup.
27 2. Loading a Control Record into the index table and obtaining the
28 associating index table references.
29 3. Obtaining the disk address of any video frame given the reference
30 and frame number within the frame range (including the ending "next"
31 frame address).

1 4. Deleting video frames from a referenced index table region.

2 5. Deleting a referenced index table region.

4 The software interface to the index table used during recording is more
5 complex. Since the index table memory region is managed internally, the
6 interface provides mechanisms to let the recording process know if the index
7 table is out of free space, or if an index table region boundary will require an
8 Available Record to be fragmented into two or more Media Records. The
9 interface for recording includes:

11 1. "Opening" free index table space for use during the recording process.

12 2. Appending frame addresses to the index table sequentially during
13 recording.

14 3. Overwriting previously stored frame addresses (used during loop
15 recording).

16 4. Removing previously stored frame addresses (used for video tape
17 emulation and frame buffer transfer features).

18 5. Indicating the free space available in the index table region for
19 recorded video frame addresses.

20 6. Freeing used index table space when aborting a recording in progress.

21 7. "Closing" index table space at the end of the recording process.

22 8. Obtaining references for associating index table regions to Records.

23 9. Writing a Control Record to a disk drive from the index table.

25 LOOP RECORDING

26 In order to record video frames to disk, the recording process of the inventive
27 circuit card needs to have both unused space on disk (i.e., one or more Available
28 Records) and unused (free) space in the index table uniquely allocated for its
29 use. This space must be uniquely allocated to the recording process to prevent
30 other processes from concurrently writing to these spaces, creating a memory
31 conflict and corrupting the recorded data. While, conceptually, all the free space

1 needed or available could be allocated just prior to recording, the inventive
2 firmware allows unused space to be allocated both prior to initiating recording
3 (Available Records are pre-allocated and become "Reserved Records") or just
4 prior to being needed during the recording process. This latter case allows other
5 processes to obtain disk space for use while recording is in progress.

6

7 The invention manages and allocates disk space by maintaining in memory a
8 directory of records on disk. The disk record directory contains the following for
9 each record on disk:

10

11 1. The disk drive unit number containing the disk record.
12 2. The disk block address of the start of the disk record.
13 3. The disk block size of the disk record.
14 4. The type of the disk record (e.g., Available Record, Media Record,
15 Control Record).

16

17 During recording, the disk record directory also contains the following for each
18 allocated Available Record being recorded to:

19

20 1. The allocation flag.
21 2. The reference to the next Available Record in the recording allocation
22 sequence.
23 3. The reference to the index table region associated with the video frames
24 recorded into the Available Record.
25 4. The number of video frames recorded into the Available Record and
26 represented in the associated index table region.

27

28 The recording process allocation firmware allocates and links in additional
29 Available Records prior to their use by filling in the reference to the Available
30 Record just allocated into the directory entry of the Available Record into which
31 video data is currently being recorded. This continues until there are no more

1 Available Records left, at which time the linkage is either terminated (loop record
2 disabled) or referenced back to the first Available Record directory entry (loop
3 record enabled). As recording into an Available Record is initiated, the reference
4 for the start of the index table region to be used is placed into the Available
5 Record's directory entry. The recording process then updates the number of
6 frames in the Available Record's directory entry, as video frames are stored on
7 disk. Note that if a recording does not use an allocated Available Record, it is
8 deallocated when recording is terminated.

9

10 As mentioned previously, after a recording is completed, the disk record
11 directory also contains the following for each resulting Control Record and Media
12 Record in the clip:

13

14 The unique clip identification number (Clip ID) for the clip.

15 The Sequence Number of the disk record within the clip.

16

17 The disk record directory thus provides the information the recording process
18 needs to seamlessly record sequential video frames in proper order to disk
19 across multiple Available Records. It also provides the information the recording
20 process needs to "loop" video frame recording when it reaches the end of
21 available disk space and overwrite previously recorded video frames in a
22 continuous manner. FIG. 4 shows an example of the relationships between disk
23 record directory entries, Available Records and index table regions for a loop
24 across three Available Records.

25

26 The uniqueness of the invention's firmware comes with the addition of the
27 requirement to loop within the allocated index table space as well as within the
28 allocated disk space. Fundamentally, a loop condition is created when either the
29 available disk space or the available index table space is exhausted. Looping
30 then overwrites frames by either overwriting the actual video frames on disk, or
31 by overwriting the frame disk block addresses in the index table, or both. In

1 addition, the overwrite process must maintain synchronization between the loss
2 of video frames on disk and the loss of their disk block addresses (“indices”) in
3 the index table. Otherwise, a condition would exist where some portion of the
4 index table contains disk block addresses that do not align to the start of
5 compressed video frames on disk. This then would cause random access
6 playback features to fail since the compressed video frame format is no longer
7 aligned to be properly decompressed. FIG. 5 shows an example where a loop
8 has just occurred. Here, a larger set of four video frames has overwritten a
9 smaller set of seven video frames. Although the indices for frames 5 through 7
10 have not been overwritten in the index table, since these frames have been
11 overwritten on disk, their indices must be deallocated from the index table as
12 well.

13

14 The invention's recording process deallocation firmware manages the
15 deallocation of both video frames on disk and their indices in the index table as
16 they are overwritten during recording. The deallocation process also currently
17 employs a simplification to the general case by looping both the video frames on
18 disk and their indices in the index table at the same time when either the space
19 on disk or space in the index table is exhausted. This causes the loop boundary
20 to always exist at the start of the space on disk and start of space in the index
21 table concurrently. This avoids the necessity in the general case of joining index
22 table fragments at the loop boundary.

23

24 Since the recording process places information regarding the video frames
25 currently being recorded into the Available Record's directory entry, the
26 deallocation firmware maintains a copy of the Available Record directory entry for
27 the deallocation point. This copy is called the “loop remnant” directory entry
28 because it refers to the remnant of the oldest video that still remains in the
29 Available Record being overwritten (see FIG. 5). The deallocation firmware
30 updates the disk block address, index table reference and number of frames in
31 the loop remnant directory entry as video frames and indices are deallocated.

1 (Note that the disk block address of the Available Record in its entry in the
2 directory is never changed, since this would corrupt the directory.) When
3 recording is finally terminated, the loop remnant directory entry represents the
4 start of the video loop (i.e., refers to the oldest sequence of video frames). If the
5 overwrite point is not on an Available Record boundary, the Available Record
6 containing the overwrite point is split into two disk records, the first being the
7 newest portion of video frames (represented by the Available Record directory
8 entry with a smaller size), and the second being the oldest portion of video
9 frames (represented by the loop remnant directory entry copied into the disk
10 record directory). The associated portions of the index table are likewise split to
11 become two Control Records. The portions of the Available Records used by the
12 recording are converted into Media Records.

13

14 With the time delay mode enabled, the recording board sends video frame
15 information to the playback board via a serial communications link. This
16 information consists of two parts, the disk address of the individual video frame
17 and information that identifies the Available Record directory entry associated
18 with the video frame location. This information allows the playback board to
19 create its own index table for the newly recorded material and generate the data
20 structures required for video playback. Like the recording board, the playback
21 board also has to deallocate video frames from the playback data structures and
22 indices from the index table as loop recording overwrites video frames. Since the
23 playback board can access any of the newly recorded video frames on disk, the
24 playback deallocation firmware actually deallocates a number of frames ahead of
25 the actual overwrite point. This prevents the playback board from accessing
26 information that is in the process of being overwritten. Also, since there is some
27 delay associated with the transmission of video frame information across the
28 serial communications link, the playback board is able to access video frames up
29 to some number of frames behind the current recording position. The playback
30 firmware accounts for the condition where playback encounters these resulting
31 endpoints by keeping the playback position within existing video through properly

1 advancing the playback frame position depending on the playback mode (e.g.,
2 stopped or playback at a particular frame rate). This allows the playback board
3 to continue playback indefinitely without interruption.

4

5 **MEMORY CONTROLLER**

6 It is common in digital systems to have a requirement wherein several
7 processes need access to a common memory. A typical means of solving such
8 a requirement is illustrated in FIG. 6. A single ported commercial memory
9 device, or bank of such devices is provided with several access paths, or
10 externally implemented ports, by which data from several sources can flow in and
11 out. Control logic must arbitrate between the processes requiring access to the
12 memory, and select a single port for connection to the memory bank for each
13 access cycle.

14

15 Quite often, the mechanism by which access to the memory is allocated is
16 either performed on a demand basis, or according to a fixed time-plan. With the
17 demand method, the process requiring usage of the memory generates a
18 "request" signal, to which the control logic responds with a "grant" when the port
19 is to be given access. With the time-plan method, access is granted at fixed time
20 intervals on a repetitive periodic basis. These mechanisms are very suitable for
21 certain types of processes which either generate memory requests in an
22 unpredictable fashion, such as random CPU accesses, or in a very regular,
23 completely predictable pattern, as with screen refresh in graphics subsystems.

24

25 The present invention implements a new method for allocating shared
26 memory bandwidth, which is particularly efficient at managing the relatively
27 unpredictable high-bandwidth data flows found in variable-rate digital video
28 streams. This mechanism can be considered as a variation of the demand-
29 driven method, and operates by making measurements of each data flow in
30 progress. A dynamic estimate is thereby developed, which is used to predict the
31 urgency of each port's requirements. This estimate is fed to the arbiter as a

1 multi-level request. The arbiter then uses the estimates to select one of the
2 contending ports for access on the next memory cycle. This procedure allows
3 the allocation pattern of memory cycles to dynamically adjust to unpredictable
4 data-flow requirements

5

6 FIG. 7 shows a 4-port memory system where multi-level arbitration is used on
7 two of the ports, numbered 1 and 2. This is similar to the configuration of the
8 video recorder board. Data flowing to or from the two ports in question is first
9 directed through staging First-In-First-Out (FIFO) data buffers connected to the
10 memory bus.

11

12 Let us first consider the case where the video data is flowing into the memory.
13 As the data in the FIFO accumulates, a multi-level request is presented to the
14 arbitration logic. This request is derived from the data level in the FIFO and
15 therefore reflects an estimate as to how soon the FIFO might be expected to
16 overflow, thus resulting in data loss and corruption of the video stream. If no
17 data is in the FIFO, the request level is "0", for no request. If some data is in the
18 FIFO, but the level is below a minimum (marked as threshold 1 in the figure), the
19 request level is "1", or low priority. In this case, the arbiter might service the port
20 only when no other is requesting service. If the data level in the FIFO is between
21 threshold levels 1 and 2, a mid-priority level "2" request is presented; whereas a
22 data level above threshold 2 results in a high-priority level "3" request. In this last
23 case, the arbitration logic could service the port on the next memory cycle
24 regardless of requests from other ports.

25

26 If the data flow is reversed such that data is flowing out of the memory into
27 the FIFO, and then to the video hardware, then the request priorities associated
28 with the FIFO data levels are reversed as well. Level "0" would then correspond
29 to a full FIFO whereas a nearly empty FIFO, with data level below threshold 1,
30 would generate a level "3" request.

31

1 FIG. 8 is a block diagram of memory control logic associated with the system
2 of FIG. 7. The arbitration block receives the requests from the ports. In the case
3 of ports 1 and 2, these requests will be multi-level. It then generates the port
4 selection signals to connect whichever port it selects for the next memory cycle
5 to the memory bus, and also a start signal to the cycle sequencing block. The
6 sequencing block generates the signals needed to control the memory chips, and
7 sends the EOCYC (end of cycle) signal back to the arbiter to enable port select.

8

9 In FIG. 9, the implementation of a simplified arbiter is shown. This logic will
10 select one of 3 requesting ports of which ports 1 and 2 have multi-level requests.
11 The multi-level requests are signaled by 3 lines for each port, P1RQL, P1RQM,
12 and P1RQH, or P2RQL, P2RQM, and P2RQH respectively for port 1 and port 2.
13 The port request logic drives only one of the 3 lines active at a time: RQL for a
14 level 1 request, RQM for level 2, and RQH for level 3. Port 3 has only a single
15 request line, P3REQ. The arbitration logic will select the highest priority
16 requesting port according to the following fixed precedence, highest to lowest:
17 P1RQH>P2RQH>P3REQ>P1RQM>P2RQM>P1RQL>P2RQL. It should be
18 noted that the actual arbitration and sequencing logic provides additional features
19 to help optimize system performance. These include arbitration for 5 ports, cycle
20 sequencing for burst transfers, and arbitration fairness. This last feature
21 prevents a single port from taking all successive memory cycles, and places an
22 upper limit on latency to grant memory service to a requesting port.

23

24 The following illustrates how multi-level port request logic, selection logic, and
25 the cycle sequencing, work together to implement dynamic arbitration in the
26 invention:

27

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46

Requesting Port

Arbitration Logic

Sequencing Logic

1. Port logic examines data flow and develops an estimate of the urgency for port service. It presents a multi-level request to arbitration representing this level of urgency.

2. Arbitration examines requests from all ports, and selects the one which appears to have the highest urgency level. Port selection signals are activated, and start of cycle is signaled to Cycle Sequencing.

4. The selected port transfers data to/from memory under control of the cycle sequencing logic.

6. Arbitration selects a port for the next cycle.

3. Sequencing generates control signals to cycle memory and transfer data.

5. The data transfer is completed and the cycle is terminated.

Dynamic Metering of Compressed Video Data Rate to Maximum

Disk Bandwidth

The user configures the invention to record video at a specified compression ratio. Two factors influence the level of compression the user may desire: 1) image quality is to a point inversely related to the level of compression, lower compression resulting in a higher quality image; and 2) maximum recording duration is directly proportional to the level of compression. In many applications maximum image quality is the primary concern. In these cases, the user typically configures the recorder for the minimum possible compression.

1
2 Any specified compression level translates directly into a target data rate
3 since incoming video is digitized to a fixed resolution, say 720 by 486 pixels at 2
4 bytes per pixel. The North American standard for motion video runs at 29.97
5 frames per second, so this yields an uncompressed data rate of $720 \times 486 \times 2 \times$
6 $29.97 = 20,974,204.8$ bytes per second. A specified compression ratio of 4:1, for
7 example, translates to a target data rate of $(20,974,204.8 / 4)$, or about 5
8 megabytes per second.

9
10 In JPEG (Joint Photographic Experts Group) compression, the method used
11 in the inventive recorder, one particular parameter controls the level of
12 compression. This parameter is the amount of *quantization* the compression
13 encoder applies to the pixel values during the succession of transformations and
14 encoding schemes specified in the JPEG compression method.

15
16 Maintaining a constant output data rate while compressing a series of
17 changing images is problematic because it is impossible to predict the results
18 from a given level of quantization: The performance depends on the image. The
19 same quantization parameters produce compressed images that vary in size by
20 an order of magnitude between, for example, a black frame and a complex,
21 highly detailed image.

22
23 Fortunately, the compression encoder used in the invention's circuit board
24 embodies a bit-rate control mechanism that varies the quantization parameters
25 automatically on an image by image basis, statistically monitoring the results to
26 converge on a specified target data rate. Since motion video usually contains
27 long series of similar images, this method is effective over an extended period of
28 time. A rapid change in the image due to a quick event in the program or editing,
29 however, generates an instantaneous peak in the size of compressed frames
30 and consequently in the short-term data rate.

31

1 The installed disk drive(s) may be unable to write the compressed data
2 stream without interruption in the event the user has configured a data rate
3 beyond the disk's capability or in the event of an extended peak in the
4 compressed data rate for the reasons outlined above. In demanding
5 applications, this results in a catastrophic failure: The only way the recorder can
6 recover is to skip frames, losing them irretrievably.

7
8 To overcome this problem, the invention's firmware includes an algorithm that
9 statistically monitors the performance of the disk compared to the video frame
10 rate to determine the disk's actual maximum data rate in the context of current
11 conditions. The algorithm also checks the state of the first-in, first-out (FIFO)
12 memory buffer between the compression encoder and the disk to insure the
13 buffer is not in danger of overflowing. If the disk is unable to sustain the user-
14 configured data rate, the firmware reprograms the compression encoder's target
15 data rate to the maximum data rate the disk can support. The firmware repeats
16 this process continuously, dynamically adjusting for changing video data rate and
17 disk performance. As conditions change, the data rate converges on the user-
18 specified compression level or the maximum data rate the disk can sustain over
19 time, whichever is less. Most important, this algorithm prevents the recorder from
20 the unacceptable deficiency of dropping frames during recording.

21
22 The flow chart of FIG. 10 details this dynamic metering algorithm as
23 implemented in the inventive recorder.

24
25 Having thus disclosed a preferred embodiment of the invention, it being
26 understood that numerous modifications and additions are contemplated and that
27 the scope of protection afforded hereby is to be limited only by the appended
28 claims and their equivalents, what is claimed is: